

Scishare: A Secure P2P Information Sharing Tool

Karlo Berket,
Abdelilah Essiari, Artur Muratas
Lawrence Berkeley National Lab

Outline



- Requirements
- System architecture
- Meeting the security requirements
 - PKI-based peer-to-peer security
- Software implementation

Requirements



- Share local data
 - Keeps data in owner's hands
 - Allows faster access to updates
 - Don't have to wait for transfer to repository
 - Central repository is no longer a requirement
 - Better scalability and fault-tolerance of system
 - Easier to consolidate existing information stores
 - Easy-to-use fine-grained access control interface is a must

Requirements



- Allow for extensible search
 - Search important feature for information systems
 - Where would the web be without search engines
 - Extensibility important
 - Goal is not to create a data representation standard
 - Allow different disciplines/applications to use the query language and data representation they are familiar with
 - Don't require translation of existing information to use the system

Requirements



- Security
 - Confidentiality and integrity of communication
 - Fine-grained access control to resources
- Support ad hoc collaborations
 - Meetings at conferences
 - Requires flexible security model
 - Quick and easy startup
 - Trust building

Requirements



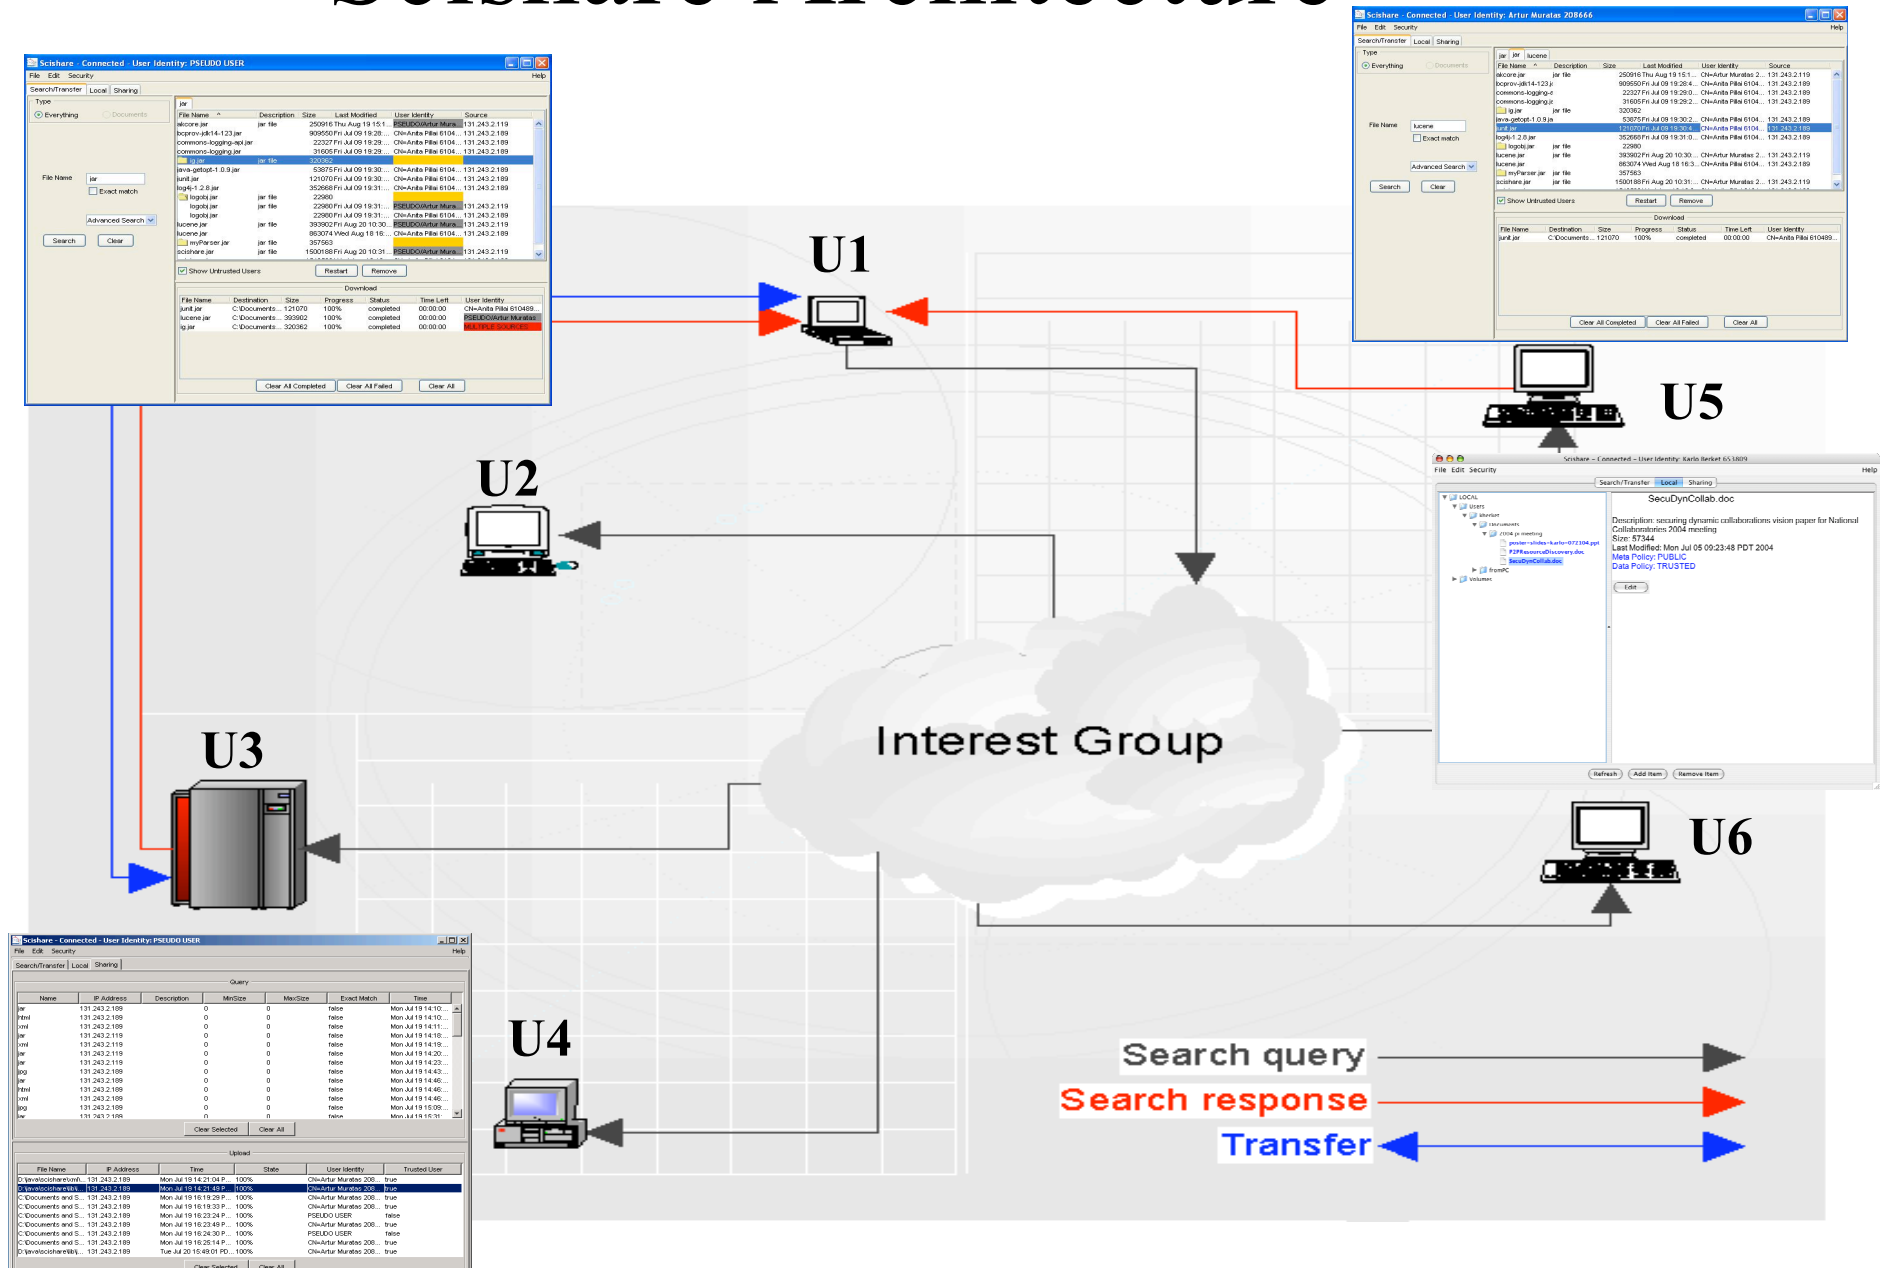
- Reach a wide community
 - Run on many OS and architectures
- Do it all in short time with limited personnel
 - Need to use existing technology whenever possible

Outline



- Requirements
- System architecture
- Meeting the security requirements
 - PKI-based peer-to-peer security
- Software implementation

Scishare Architecture



Outline



- Requirements
- System architecture
- Meeting the security requirements
 - PKI-based peer-to-peer security
- Software implementation

Security Goals



- Confidentiality and integrity of communication
- Fine-grained access control to resources
- Support ad hoc collaborations
- Assumption
 - X.509 identity certificates

Approach



- Use Public Key Infrastructure (PKI)
 - X509 certification/online CAs
 - Flexible Trust Models
 - Reduces Key Management issues
- Use existing PKI-based security technologies
 - Modifications are external
 - Reduce the risk of introducing security holes

Traditional Security Model



- Authorized users are predefined
 - In or out (of system)
 - Harder to meet ‘new people’ online in a collaboration
- Policies are managed by third party entities (administrators)
 - Valid in many use cases
 - Hard to start a spontaneous collaboration
 - Setup takes time
 - Hard to invite a person to an established collaboration
 - Must contact resource administrators
- Security becomes a nuisance
 - Users may resort to insecure solutions

A Flexible Security Model



- Partition the collaboration into two types of **secure** components:
 - Public
 - Capture users' identities
 - Gradual trust in the collaboration
 - Turn off public components => traditional model
 - Protected
 - Authorized users only
 - Give invitation/escort powers to some of these users
- Example of components:
 - Communication channels, online instruments, chat rooms, shared spaces, files, ...

Components in scishare



- Unicast channels
 - Managed by the users participating in the communication
- Multicast channel
 - Managed by ‘Third-Party administrators’
- Files and metadata
 - Managed by individual users

Securing Unicast



- Secure communication channel (SSL)
 - Confidentiality
 - Integrity
 - Authentication
 - Typically only server presents X509 certificate
 - Require both parties to present X509 certificate (mutual authentication)
 - Every user needs a certificate

Securing Unicast



- Provide users with pseudo (self-signed) X509 certificates if they don't have any
- Custom trust manager
 - Accepts any valid chain
 - Marks users as trusted if user and chain verify
 - Remember un-trusted users
 - Can later authorize un-trusted users based on experience
- A single channel can handle both protected and public traffic
 - Simplifies development

Securing Multicast



- Need a secure group communication channel with properties similar to SSL:
 - implements an authenticated and encrypted group channel
 - enables group members to establish a session key
 - certificate-based access control

Securing Multicast



- Public group communication channel
 - Every user can join
- Protected group communication channel
 - Fine-grained access control
 - Join, invite, escort
 - Capabilities
 - Short lived, signed by the enforcers
 - Invitations/Escorts
 - Short lived, signed by authorized users
- A single communication channel
 - A protected channel over a public one (sub-group)

Securing Data and Metadata



- Provide a simple high level interface to users
 - Manage policies
 - Manage groups
- Authorization engine is used underneath (e.g. Akenti)
 - Distributed groups
 - User revocation
 - Future complex expressions
 - Time of day, ...

Read More About It



K. Berket, A.Essiari and A. Muratas

PKI-Based Security for Peer-to-Peer Information Sharing

Proceedings of the Fourth IEEE
International Conference on Peer-to-
Peer Computing, Zurich, Switzerland,
Aug. 25-27, 2004.

Outline



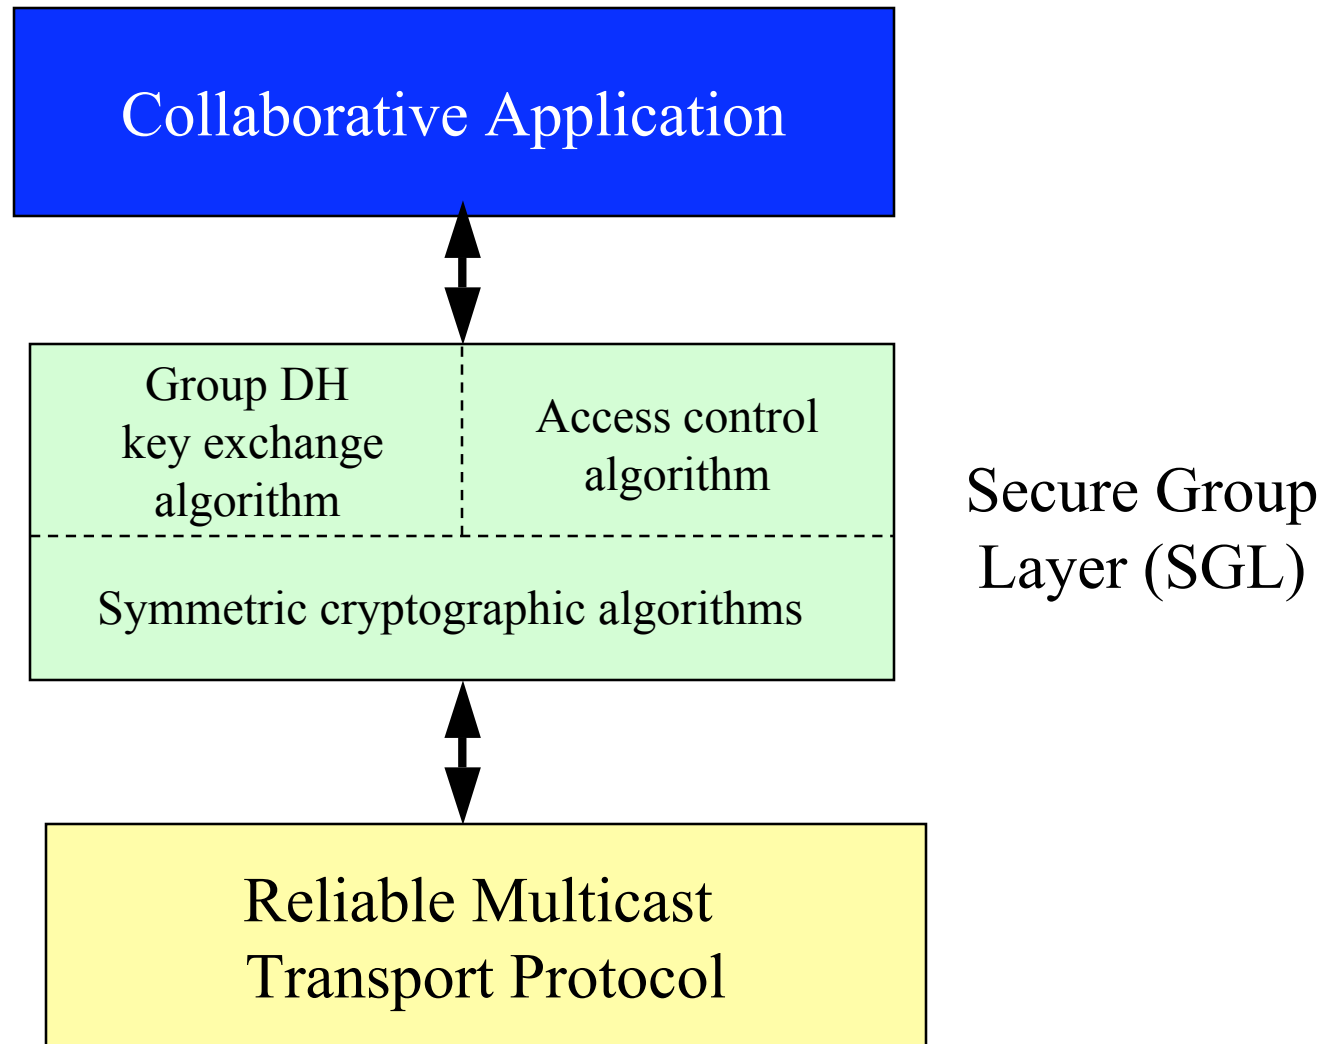
- Requirements
- System architecture
- Meeting the security requirements
 - PKI-based peer-to-peer security
- Software implementation

Build on Existing Tools



- XML messaging
 - JAXB to generate code from schema
- Securing the group communication channel
 - Secure Group Layer (SGL)
 - InterGroup
- Authorization engine
 - Akenti

Security at the Multicast-Transport Layer



The Reliable Multicast Transport Layer



- Provide SGL with reliable and ordered delivery of messages
 - data messages are delivered in order - FIFO, partial, and total - at each member of the group
- Provide SGL with membership notifications
 - membership changes delivered in order with respect to data messages
- Several systems provide a reliable multicast layer
 - e.g., Totem and InterGroup

The Secure Group Layer




- Symmetric crypto algorithms
 - implement an authenticated and encrypted channel
- A group key-exchange cryptographic primitive enables group members to establish a session key
- A certificate-based access control mechanism makes sure that only the legitimate parties have access to the session key
 - off-line (does not participate in key exchange)

Group Key-Exchange



- Up-flow: U_i raises received values to the power of x_i and forwards to U_{i+1}
- Down-flow: U_n processes the last up-flow and broadcasts

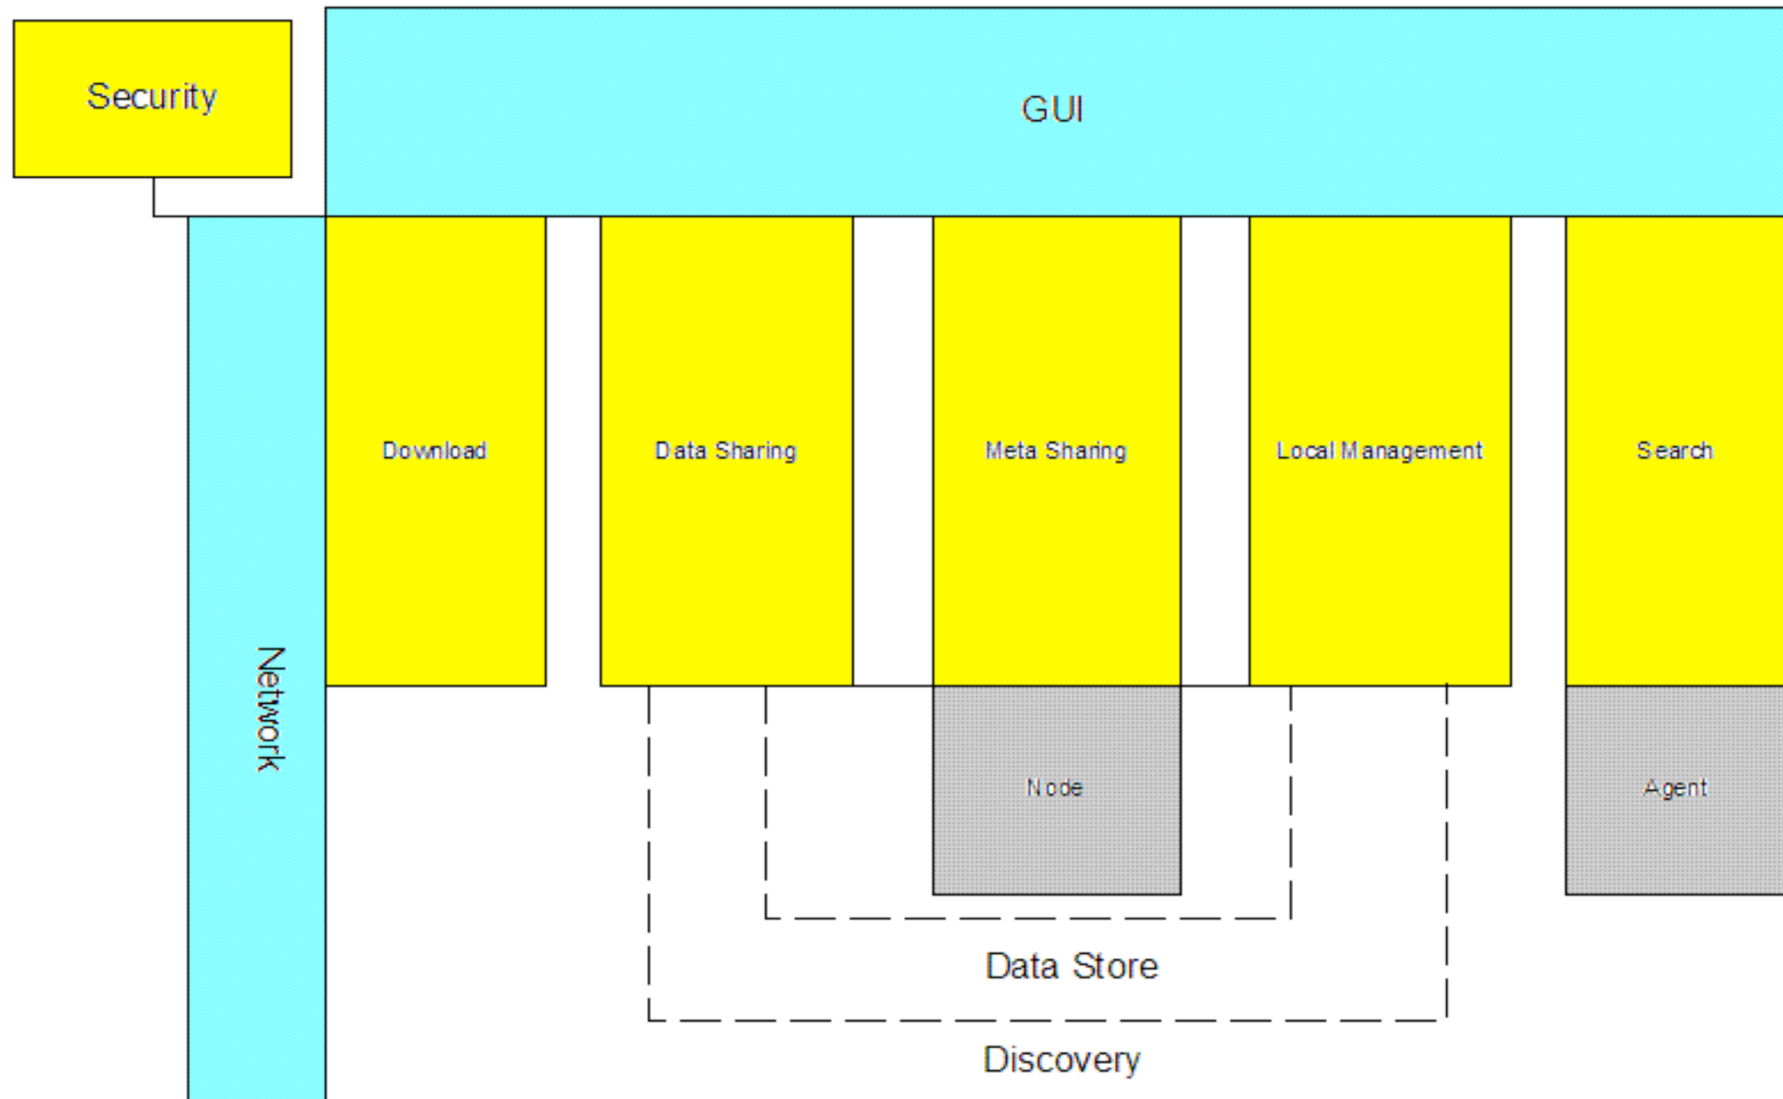


Akenti Distributed Authorization



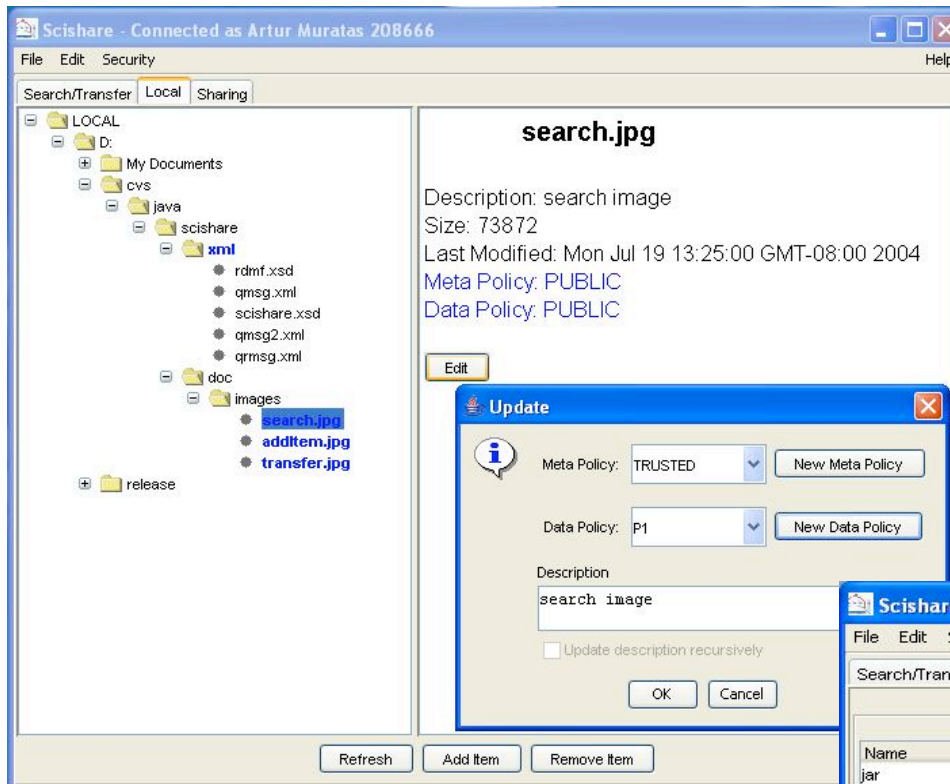
- Target widely distributed environments
 - Resources (instruments, executables, ...)
 - Principals:
 - Resource owners (stakeholders)
 - User-Attribute Issuers
 - Users
- Collaborative/Grid environments that could span many autonomous/dispersed organizations.
- Provide a flexible and secure way for stakeholders to remotely and independently define authorization policy and allow fine-grained access control.

Software Architecture



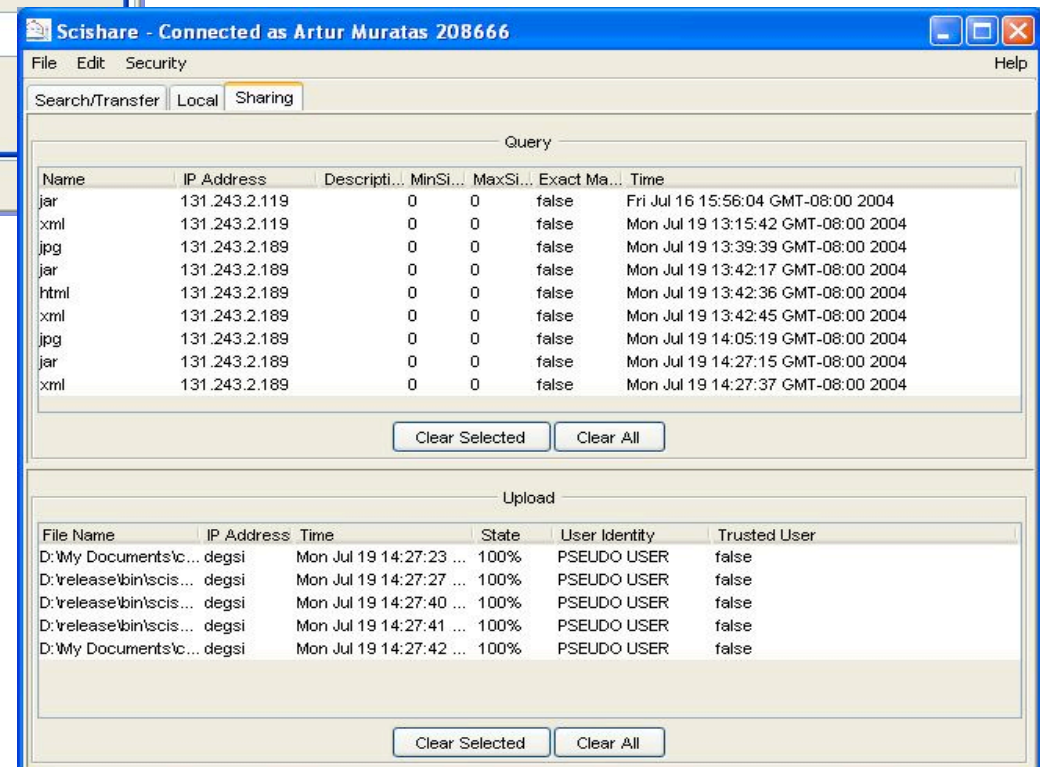
Local

- Add/remove files in DB
- Synchronize DB with file system
- Map different policies to metadata and data



Sharing

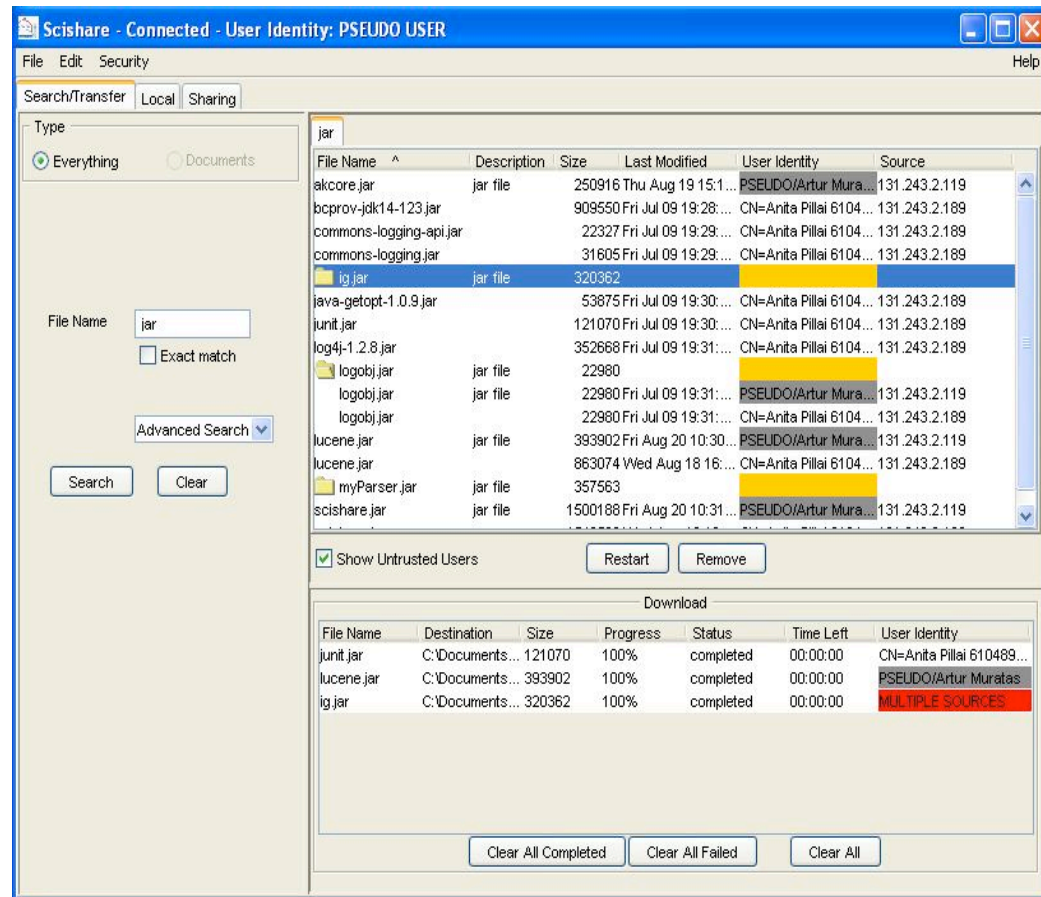
- View incoming queries
- View file uploads




Search - Transfer



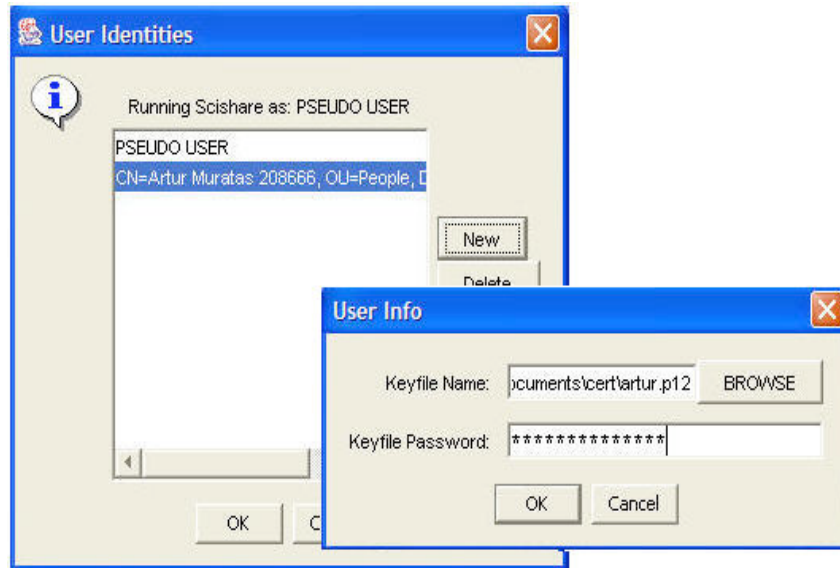
- Create basic and advanced queries
- Start a search
- Group search results based on the same hash
- Display the origin of the metadata and its trustworthiness
- Allow user to download portions of a file in parallel
- Display the origin of the file and its trustworthiness




Manage CAs




Manage User Identity



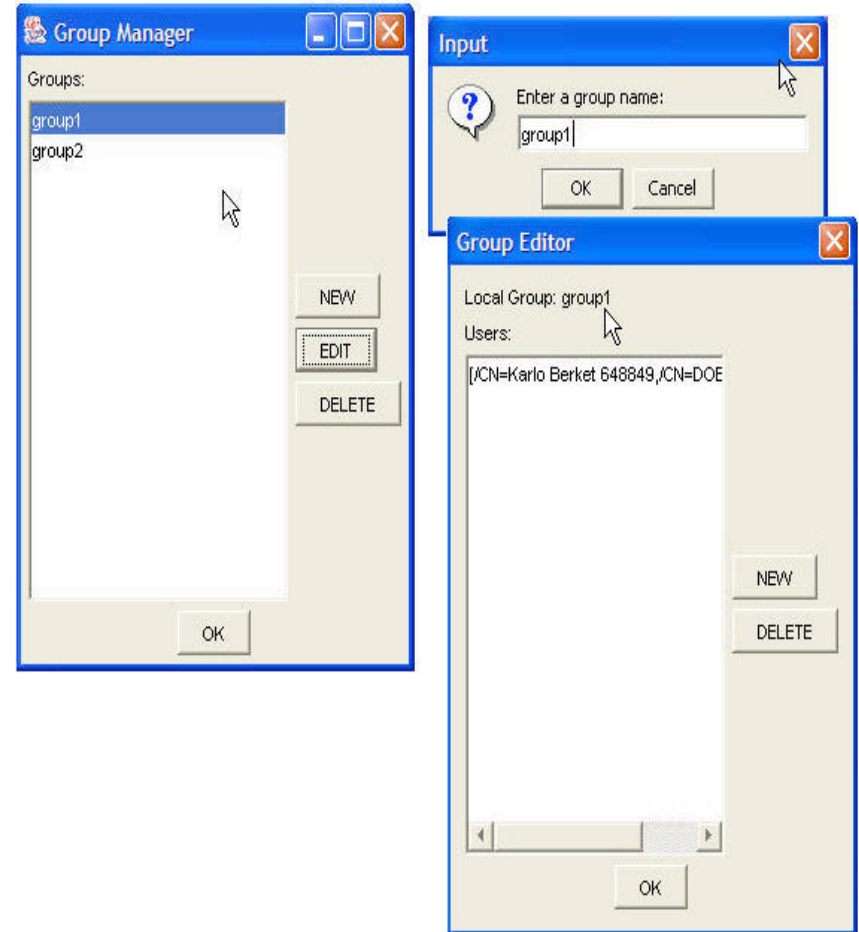
View Un-trusted Users



Manage Policies



Manage Groups



On-going Work



- Performance measurements (security impact, scalability, etc.)
- Remote groups
- Securing the queries
- Providing access control to search group

More Information

- Project page:
<http://www.dsd.lbl.gov/P2P/file-share>
- Software download:
<http://www.dsd.lbl.gov/scishare>
- E-mail: kberket@lbl.gov